Explaining Choice Behavior: The Intentional Selection Assumption

Kelley Durkin[1], Leyla R. Caglar[2], Elizabeth Bonawitz[2], & Patrick Shafto[1]
(kelley.durkin@gmail.com; lcaglar@psychology.rutgers.edu; elizabeth.bonawitz@rutgers.edu; patrick.shafto@gmail.com)
[1] Department of Mathematics & Computer Science; [2] Department of Psychology
Rutgers University – Newark; 101 Warren Street, Smith Hall; Newark, NJ 07102 USA

Abstract

How do people decide between several options presented to them? Normative accounts suggest the utilities of options are fixed, but subjective accounts suggest utilities depend on context. In the current paper, we propose a novel model of choice that may help reconcile these accounts. We propose that choice behavior may depend on an "Intentional Selection Assumption": when people are presented with multiple options, they assume the options were intentionally selected by a person with specific questions in mind. Inferences about the intentional selection of options inform the chooser about the features that are intended to be most relevant. In this way, context can affect the desirability of a particular option, without requiring shifting utilities over features. Two behavioral experiments support the claim that participants are sensitive to intentional selection. We discuss the importance of taking choosers' assumptions about intentional selection into account in future investigations of choice behavior.

Keywords: intentional selection; choice; computational modeling

When people are presented with multiple options, how do they choose among the possibilities? For example, consider a choice about which newspaper medium to purchase: an online subscription for $59 or a joint online and print subscription for $125. When presented with these two subscription options, people tend to choose the online subscription. However, if a third print-only subscription option is added for $125, people tend to choose the joint online and print subscription for the same price (Ariely, 2010). What may explain this sudden shift in preferring the online and print $125 option? There are many features you could consider. You may consider which subscription is the cheapest or most convenient. However, if you know that the newspaper is presenting what they believe are all useful and valuable subscription options, then you may impute different meaning to the options. The newspaper's intentional choice to present multiple options that include print access may affect one's choice and the perceived utility of print access.

The question of how people choose from a set of options underpins a wide variety of disciplines including psychology, economics, marketing, and computer science; this interdisciplinary interest has resulted in a variety of formal models of choice behavior (Luce, 1959; McFadden, 1977; Sutton & Barto, 1998; Thurstone, 1927; Yellot, 1977). Choice tasks also drive behavioral experiments and survey research, implying a critical, but often overlooked, role in linking theoretical questions to empirical data.

To explain choice behavior, Duncan Luce (1959) proposed the Luce choice rule, which captures the notion that choice is systematically probabilistic and normative (i.e., options maintain fixed utility). The Luce choice rule suggests that the probability of each item is proportional to its utility relative to the other presently available items; items are selected in proportion to their weight.

However, empirical research has cast doubt on this normative account (e.g. McFadden, 1977; Simonson & Tversky, 1992; Tversky, 1972). For instance, consider our newspaper subscription example. When only presented with two options (the online and the joint print and online subscriptions), people presumably saw higher utility in the online subscription and chose it more often. However, adding a print subscription at the same price as the joint print and online subscription presumably increased the perceived utility of print access, such that more people start to choose the more expensive joint print and online option. Based on this and numerous other phenomena, researchers have argued that choice is not normative, but that utility judgments are idiosyncratically affected by context and are therefore not stable.

Subjective accounts of choice create a different set of challenges. For example, they might suggest that people do not have a stable concept of utility at all. However, without a stable concept, it is unclear how to explain the degree of systematicity observed in choice behavior. Thus, it remains an open question how to explain the systematic variability of choices, especially as a consequence of context.

Here we propose a novel factor, the "Intentional Selection Assumption", which may influence people's choice behavior. We suggest that when people are provided with a set of options to choose among, they treat the set of options as intentionally selected by a person with a specific question in mind. That is, choosers may consider the goals and beliefs of the individual presenting the options in order to help resolve uncertainty about the relative importance of the options' features. This proposal provides a novel application of recent research formalizing learning from others (Shafto, Goodman, & Frank, 2012). It also provides new empirical predictions pertaining to classic results in choice.

In what follows, we first detail how these social assumptions shape choice behavior. We then present two novel experiments that manipulate whether choice options are provided intentionally or randomly. Our results demonstrate that adults are sensitive to intentional selection, providing support for the Intentional Selection Assumption.

We conclude with a discussion of how our model may provide a middle ground between past accounts of choice.

Intentional Selection Assumption

Recent research has investigated the effects of intentional selection on learning. This research suggests that learners attend to what examples are and are not presented because they assume examples are selected by another person for a reason (Bonawitz, Shafto et al., 2011; Shafto & Goodman, 2008; Shafto, Goodman, & Griffiths, 2014). Shafto, Goodman, and Frank (2012) have proposed a framework for formalizing these sorts of social effects on learning. Bonawitz, Shafto et al. (2011) have applied this approach to explain children's exploratory play. (See also Goodman, Baker, & Tenenbaum (2009) on casual inference, Shafto, Eaves, Navarro, and Perfors (2012) on epistemic trust, and Frank and Goodman (2012) on communication). An important contribution of this work is to focus on the inferential affordances provided to the learner by leveraging intuitive psychological reasoning. Because people's actions are goal-directed, rather than random, we can reason about why they do things, and this has implications for the kind and strength of inferences that can be drawn.

However, current theories of choice do not factor in the importance of the intentional selection of examples. They either do not mention the process by which options are selected or explicitly assume that the presented options are randomly sampled. Indeed, many models directly or indirectly assume independence of irrelevant alternatives (IIA) (Arrow, 1963). IIA states that eliminating some alternatives should not affect the individual preferences between remaining options. In other words, if option A is preferred to option B, introducing a third option, C, should not make B preferable to A. In this way, alternatives are irrelevant to the choice between two options.

Empirical evidence suggests that human choice behavior is not independent of the alternative options (e.g., Ariely, 2010; Simonson & Tversky, 1992, Tversky, 1972). One classic example, related to the newspaper example presented here, is known as the Compromise Effect (Simonson & Tversky, 1992). In this task, participants chose between cameras that varied in price and quality. When participants chose between a high price, high quality camera and a low price, low quality camera, each camera was selected about the same number of times. When a third camera was added to the set of options at either extreme (e.g., lower quality and lower price), then the intermediate camera was chosen more frequently than the previously equally favored one. This effect of context on people's choice behavior illustrates that choice depends on more than normative weights of options, which are assumed to be independent of each other.

We suggest that in these previous tasks, people may not have assumed that the options were randomly sampled. Previous experiments were often ambiguous about how the options presented to choosers were selected. Participants may have instead inferred that options were intentionally selected by a person with a goal in mind. In particular, participants may believe that those particular options were chosen for them because the presenter wanted to stress the relative importance of particular features. If so, it would suggest that the Intentional Selection Assumption may also provide an explanation for people's seemingly subjective choice behavior.

Intentional Selection Model

Any set of options has a potentially infinite number of features. So how might a chooser decide which features are most relevant to consider in a particular context? We know, for example, that a life raft is a better choice than a laptop if you are drowning, and a laptop would be the preferred choice for composing a paper, but options and contexts often present greater uncertainty about the relevant features to consider than this trivial example illustrates.

How might we alleviate uncertainty in choice? The Intentional Selection Assumption depends on the notion that participants are evaluating options as if those options were chosen by someone with a goal in mind. This goal could be to highlight the features that the participants should consider relevant for making a particular choice. The evidence that a participant is given to evaluate the likely goals of the option-selector (and thus evaluate the likely relevant features) is the set of options provided. Different sets of options can thus highlight different intentions of the selector, and in turn, highlight different features of relevance. In this way, a change in the set of intentionally selected options allows the chooser's utilities of those options to change. However, this does not require a subjective assumption because the feature weights are not changing from instance to instance. That is, in our approach, utilities are stable over features, but the chooser may have uncertainty about which features are relevant in the particular context. The *relevance* serves as a context-dependent weight over features. The selected options provide information as to which features are most relevant to consider in a given context. As a result, the overall utility of an item will depend on the relevant features that provide said utility, and the inferences about which features are most relevant will be context-dependent. Our approach can therefore be seen as a middle ground between the normative and subjective approaches.

A formal model capturing these notions is given in Shafto and Bonawitz (*in press*), but we provide an intuitive account of the implications of an Intentional Selection Assumption here. Our approach begins with the idea that an option is a composition of features, such that an option's overall utility is the sum over the weighted utilities of all possible features. Our approach diverges from traditional models of choice (e.g., Restle, 1961; Tversky, 1972) in that we propose that an item's utility is determined by the utility of its *relevant* features. Therefore, the utility of an option is a function of the weights of its features, each of which may or may not be relevant in a given context. While Ariely (2000) has proposed that relevance affects an item's utility, this past

model has not provided a thorough explanation for why that may be.

An important challenge is specifying how feature relevance is assessed. In our approach, feature relevance depends on the full set of options selected, as well as an inference about the intention of the questioner in providing the examples. The chooser assumes that the observed selection of options is chosen intentionally, with a specific question in mind. The chooser, observing the selected items, can reason about the intended question and use that inference to constrain the uncertainty about which features are relevant. The chooser will have uncertainty about the intentions of the questioner. Thus, the probability of relevance must take into account many possibilities.

How might we evaluate the probability of observing a set of options given a particular hypothesis about relevance? Intuitively, we might believe that in order to discern relevant features, the ideal set of options would contrast in utilities among those features. For instance, if the feature red was relevant, to highlight this fact, a questioner would prefer to select options that contrast along this feature, leading to a set containing options that are red and options that are not red. Similarly, in the case of a dimensional feature, such as price, a questioner should choose options that contrast (e.g., one high and one low) to emphasize the variability. We test this intuition in Experiment 1.

Questioners are not limited to conveying information about merely which features are relevant; the options can also convey knowledge about the distribution of the objects. If two examples are selected to contrast strongly along a dimension, the selection of a third between those extremes becomes an indicator of the middle of that distribution. Thus, a set that spans a "representative" sample of utilities over the features of relevance provides stronger support for that relevance hypothesis than a set that does not distribute evenly (Shafto et al., 2014; Tenenbaum & Griffiths, 2001). For example, if two options are provided, one that costs less and provides one service (e.g., $59 for an online subscription) and one that costs more and provides two services (e.g., $125 for a print & online subscription), it is unclear which cost feature is most relevant. However the addition of a third option (e.g., $125 for a print subscription) provides information about representativeness. This highlights the utility of a particular feature (e.g., print access), suggesting it should receive a higher relevance weight. Experiment 2 explores this prediction.

Importantly, in this proposed approach, the relevance of each feature of an option can vary depending on context in interesting ways. It depends on the full set of options selected, as well as an inference about the intention of the questioner in providing the examples. Thus, context plays a role in helping the chooser infer which features are relevant, and our model provides a quantitative way to explain past context-dependent effects of choice behavior.

Current Experiments

The current experiments aim to qualitatively illustrate that intentional selection can be used to explain behavioral data of choice. Experiment 1 examines how the intentional selection assumption affects the inferences people make about the beliefs of others. Experiment 2 examines how the intentional selection assumption affects people's choices when presented with varying contexts and features of options. In Experiment 2a and 2b, we included a baseline control to replicate Ariely's (2010) findings. In all experiments we include a condition where options were intentionally selected and a condition where options were randomly or accidentally included. By comparing the choices made when options are selected intentionally versus accidentally, we can assess whether people's choice behavior is sensitive to the intentional selection assumption.

Experiment 1

Method

The experimental participants included 95 workers from Amazon Mechanical Turk. We initially offered participants $0.10 payment for completion of the two-minute study, but then offered $0.50 pay to encourage recruitment. Twelve participants failed to pass a simple attention check and were consequently excluded from the analysis. The remaining 83 participants were randomly assigned to one of two conditions: the *random sampling* condition (N=43) or the *intentional* condition (N=40). In each condition participants were asked to make a forced choice about a novel object labeled Widget.

In the *random sampling* condition, participants were told that a coin flip determines the shape (triangle or square) and the color (blue or red) of the Widgets being produced in a factory. Participants were told that Widget 1 is triangular and blue, and that Widget 2 is square and red. Then participants were told that Widget 3 is triangular but that the text information about the color of Widget 3 was cut off. Participants were asked to guess whether the color that the machine generated for Widget 3 was blue or red.

In the *intentional* condition, participants were told that their friend (who was visiting the Widget factory) was sending a text message to ask which Widget they would like to receive as a gift. The Widget choices presented in the text were identical to the *random sampling* condition: Widget 1 (triangular and blue), Widget 2 (square and red), and for Widget 3 the text was cut off after the word "triangular." Participants were asked to guess whether the color that the friend was offering for Widget 3 was blue or red.

Results and Discussion

We first compared whether participants differed in their blue-red choices by condition (Figure 1). A chi-square revealed significant differences between conditions, χ^2 (1, N=83) = 11.62, $p < .001$. Consistent with our predictions, in the *intentional* condition, participants selected "red"

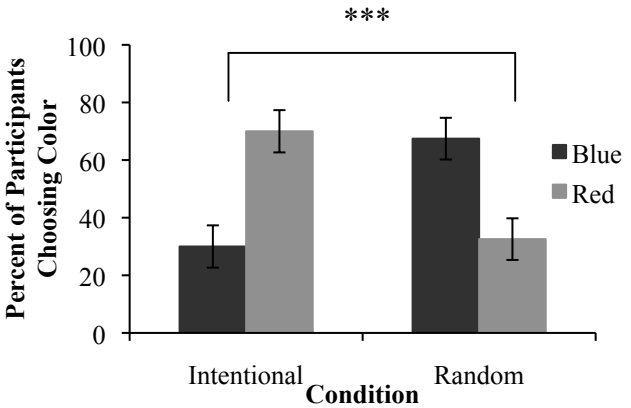

Figure 1: Exp. 1. Participants in the *intentional* condition were more likely to choose "red" while participants in the *random* condition were more likely to choose "blue". Error bars represent standard errors.

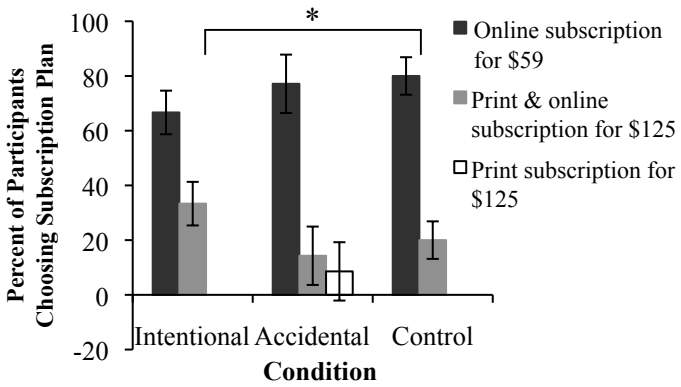

Figure 2: Exp. 2a. Participants were more likely to choose the print & online subscription option in the *intentional* condition than other conditions. Error bars represent standard errors.

significantly more than predicted by chance (N selecting red = 28 of 40; binomial, $p < .01$). Interestingly, in the *random sampling* condition participants were more likely to choose the "blue" choice than predicted by chance responding (N selecting red = 14 of 43; binomial, $p < .05$).

The bias to choose "blue" in the *random sampling* condition is consistent with previous research that suggests that adults over attribute alternation to random events (Reichenbach 1934/1939; see also Bar-Hillel & Wagenaar, 1991 for review). In our task participants were told that the machine had produced one blue object and then one red object, so the alternation account predicts the observed bias for choosing blue for the final widget in the *random sampling* condition. Importantly, our model of intentional choice predicts that participants in the *intentional* condition should believe their friend was intending to produce a unique sample of options, and thus the final widget was likely a red triangle (making it different from the other offered blue triangle and also breaking the alternation bias).

Experiment 2a

Method

112 workers from Amazon Mechanical Turk participated for $0.50 payment. Six participants were dropped because they failed to pass attention checks. The remaining 106 participants were randomly assigned to one of three conditions: the *intentional* condition (N=36), the *accidental* condition (N=35), or the *control* condition (N=35).

Participants completed an online survey based on past work on relativity of choices (see Ariely, 2010). Participants were told that they needed to choose a type of subscription plan to purchase to access a fictional newspaper. The options presented differed depending on condition. In the *intentional* condition, participants were presented descriptions of an online subscription for $59, a print subscription for $125, and a joint online and print subscription for $125 and asked to choose one to purchase. In the *accidental* condition, participants were presented with the same three subscription plans and asked to choose one. However, they were told that the website had accidentally not been updated so the print-only subscription option was not meant to be available (although they could still purchase it now if they wished). In the *control* condition, participants were presented just two options: the online subscription for $59 or the online & print subscription for $125.

Results and Discussion

The number of participants who chose each subscription plan differed overall between conditions, $\chi^2(4, N=106) = 9.56$, $p = .048$, with participants in the *intentional* condition being more likely to choose the print & online subscription than the other conditions (Figure 2). Indeed, the *intentional* and *accidental* conditions were significantly different from one another even though the same subscription plans were presented to both conditions, $\chi^2(2, N=71) = 6.05$, $p = .049$.

We used the ratio between the online choice and the joint print and online choice in the control condition to assess whether participants demonstrated similar preference in the *intentional* and *accidental* conditions. In the *control* condition, the preference of the online choice to the joint print and online choice was 4:1. In the *intentional* condition, there was a significant difference from this ratio, by binomial test against .2, $p = .042$, suggesting that participants in the *intentional* condition were using additional information that all of these options were chosen by the newspaper to evaluate the weights of the options. These results replicate past experiments of choice, which find a difference in participant responses when a third option is presented. We argue that this difference is due to an intentional selection assumption. To test this, however, we must show that when intention is removed and a third option is presented accidentally, this 'third-option-effect' goes away. Indeed, this is what we found; comparing the *accidental* condition to the ratio given by the control group revealed no differences, $p = .360$. Thus, when the intentional selection assumption was removed the classic context effects disappeared. The *accidental* condition likely chose similarly to the *control* condition because they did not see the print option as having added value even with the print-only option presented. However, in the *intentional*

condition, there was a significant difference from this ratio, p = .042, suggesting that, unlike the *accidental* condition, participants in the *intentional* condition were using additional information that all of these options were chosen by the newspaper to evaluate the weights of the options.

Given these results, presenting options as either intentionally or randomly selected changes the choices people make. In particular, these results support the intentional selection account. When participants believed the print subscription was presented to them intentionally to achieve the goals of the newspaper, this may have increased the perceived utility of having print access compared to the control condition. The perceived weight of the online subscription compared to the joint print and online subscriptions was also lower for participants in the intentional condition. However, when participants believed the print subscription was only provided accidentally, the participants chose similarly to the control condition and did not see as much utility in the print subscription.

Experiment 2b

Method

100 workers from Amazon Mechanical Turk participated. Three participants were dropped because they failed to pass attention checks. The remaining 97 participants were randomly assigned to the *intentional* condition (N=32), the *accidental* condition (N=32), or the *control* condition (N=33).

The subscription price for the online subscription was changed from \$59 to \$99. This allowed us to test the prediction that making the price feature of the options more similar would increase the differences between the *intentional* and *accidental* conditions. Also, after participants chose a subscription plan, we asked them how important it was for them to have online and print access to the newspaper on a scale from 0 (Not important) to 100 (Extremely important).

Results and Discussion

Replicating Experiment 2a, the number of participants who chose each subscription plan differed overall between conditions, $\chi^2(4, N=97) = 10.20$, p = .037, with participants in the *intentional* condition being more likely to choose the joint print and online subscription and less likely to choose the online subscription than the other conditions (Figure 3). Indeed, the *intentional* and *accidental* conditions were significantly different from one another even though the same subscription plans were presented to both conditions, $\chi^2(2, N=64) = 7.64$, p = .022.

The *control* condition provided an expected weight of preference between the online choice and the joint print and online choice (about 2:1). The *accidental* condition was not significantly different from this weight, by binomial test against .33, p = .259. However, in the *intentional* condition, there was a significant difference from this weight, p = .002, suggesting that, unlike the *accidental* condition, participants

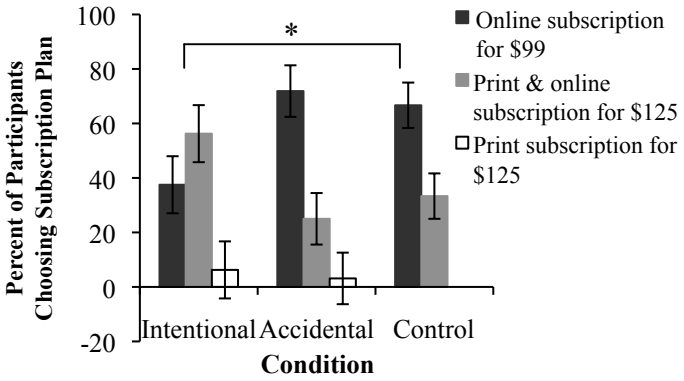

Figure 3: Exp. 2b. Participants were more likely to choose the print & online subscription option and less likely to choose the online subscription option in the *intentional* condition than other conditions. Error bars represent standard errors.

in the *intentional* condition were again using information about intention to evaluate the weights of the options.

We also compared the importance ratings of the online access and print access to the newspaper. There were no significant differences between conditions in how important online access was, $F(2, 94) = 1.09$, p = .341, η_p^2 = .023. However, there were significant differences between conditions in how important print access was (M = 51.91, SD = 38.05 for the *intentional* condition; M = 31.19, SD = 34.15 for the *accidental* condition; M = 32.18, SD = 26.50 for the *control* condition), $F(2, 94) = 3.99$, p = .022, η_p^2 =.078. Consistent with the intentional selection model, contrasts revealed that participants in the *intentional* condition rated print access as more important than the *accidental* and *control* conditions, $F(1, 94) = 6.24$, p = .014, η_p^2 = .062 and $F(1, 94) = 5.74$, p = .019, η_p^2 = .058.

These findings suggest that assuming intentional selection of options may have resulted in people making different choices than when that intention was not present. It is important to note that wanting to decrease dissonance with their previous choice could also have influenced participants' ratings. Assuming intention may have increased the perceived utility of print access compared to the accidental and control conditions. Participants may have considered the newspaper's goal to maximize profits and used that to evaluate the perceived utility of the presented options. These findings, when compared with results from Experiment 2a, also suggest that increasing the price of the online subscription option decreased the relevance of the price feature of that option, and thus it increased the difference between the intentional and accidental conditions. When price became a less relevant feature, participants possibly attended more to the fact that someone was intentionally selecting options and then inferred that it meant there was an increased utility to having print access.

General Discussion

Past models of choice have included normative approaches, which assume all options have fixed utilities, and subjective approaches, which assume utility is relative to the options presented. For both approaches, there have been many context-dependent empirical results that are difficult to

explain. One reason these effects may be hard to explain is that people may be reasoning about the intentional selection, or lack thereof, of the options, but past experiments have not explicitly investigated this possibility.

We provided a novel account, the Intentional Selection Assumption, and we are the first to account for the role of intentional selection when making choices. We suggested that participants may have a fixed notion of feature weight, but that the *relevance* of features could be context dependent. In particular, the set of options provided by an intentional agent, with a goal in mind, could provide information as to which features are most relevant to the agent. The participants can then use this information about what the agent sees as relevant combined with their knowledge about the agent's goals. If options are presented randomly, the same inferences are not warranted.

The results from the current experiments suggest that the intentional selection assumption guides people's inferences about feature relevance. Participants seem to consider the goals and knowledge of the questioner and use the selected set of options to inform this inference.

Our results showed that when participants were convinced that options were randomly selected, previously observed effects (e.g., the asymmetric dominance effect) were nullified. This has important implications for explanations of these effects. We suggest that models of choice behavior must account for the intentional selection assumption, which can explain these results. For instance, we replicated Ariely's (2010) finding that people are more likely to choose the joint print and online subscription option than the online option if they are also presented with a print subscription of similar price. However, in the current experiments, this was only true when all options were assumed to be intentionally selected. When the print-only option was presented accidentally, participants chose similarly to the control condition that only saw two options. Thus, the asymmetric dominance effect may rely on this Intentional Selection Assumption.

There is a strong link in education between learning and choice. When teachers present examples or multiple-choice options, students may make inferences about which features of concepts are important. This can affect the choices students make in learning, as well as how they create conceptual structures. Marketing and teaching will likely entail different assumptions by the learners about the goals of the selectors – with intentional deception being more likely in marketing and helpfulness being an important goal of teaching. Nonetheless, our model shows the importance of considering the goals – whatever they may be – in models of choice. Future empirical and modeling work should investigate the important role intentional selection plays in choice and learning.

The choice literature is vast, and there are many interesting effects related to choice behavior. We have developed an account based on the intentional selection assumption and conducted a first qualitative test with a few examples. However, these are only the first steps in exploring this important interpretation of choice behavior; we have presented just a few options for consideration.

Acknowledgments

This research was supported in part by funding from the National Science Foundation, CAREER grant DRL-1149116 to PS.

References

Ariely, D. (2000). Controlling the information flow: Effects on consumers' decision making and preferences. *J. of Consumer Research, 27*, 233-248.

Ariely, D. (2010). *Predictably irrational: The hidden forces that shape our decisions*. New York: HarperCollins.

Arrow, K. J. (1963). *Social choice and individual values*. New Haven, CT: Yale University Press.

Bar-Hillel, M., & Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics, 12*, 428-454.

Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Teaching limits children's spontaneous exploration and discovery. *Cognition*, *120*, 322-330.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science, 336*, 998.

Goodman, N. D., Baker, C. L., & Tenenbaum, J. B. (2009). Cause and intent: Social reasoning in causal learning. *Proceedings of the 31st ann. conf. of the Cog. Sci. Soc.* (pp.2759-2764). Hillsdale, NJ: Lawrence Erlbaum Associates.

Luce, R. D. (1959). *Individual choice behavior*. New York: Wiley.

McFadden, D. (1977). Quantal choice analysis: A survey. *Annals of Economic and Social Measurement, 5*, 363-390.

Reichenbach, H. (1934/1949). *The theory of probability.* Berkeley, CA: University of California Press.

Restle, F. (1961). *Psychology of judgment and choice*. NY: Wiley.

Shafto, P., & Bonawitz, L. (in press). Choice among intentionally selected options. In B. H. Ross (Ed.), *Psychology of Learning and Motivation* (Vol. 63), New York: Academic Press.

Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science, 15*, 436-447.

Shafto, P., & Goodman, N. D. (2008). Teaching games: Statistical sampling assumptions for pedagogical situations. *Proceedings of the 30th ann. conf. of the Cog. Sci. Soc.* (pp.1632-1637). Hillsdale, NJ: Lawrence Erlbaum Associates.

Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psych.l Sci., 7*, 341-351.

Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology, 71*, 55-89.

Simonson, I., & Tversky, A. (1992). Choice in context: Tradeoff contrast and extremeness aversion. *J. Mark. Res., 29*, 281-295.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Tenenbaum, J. B., & Griffiths, T. L. (2001). The rational basis of representativeness. *Proc. of the 23rd ann. conf. of the Cog. Sci. Soc.* (p.1036-1041). Hillsdale, NJ: Lawrence Erlbaum Associates.

Thurstone, L. L. (1927). A law of comparative judgment. *Psych. Rev., 34*, 273-286.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psych. Rev., 79*, 281-299.

Yellot, J. I. (1977). The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution. *J. of Math. Psych., 15*, 109-144.